arXiv:1511.06485v4 [cs.LG] 8 Feb 2016

# THE EFFECT OF GRADIENT NOISE ON THE ENERGY LANDSCAPE OF DEEP NETWORKS

PRATIK CHAUDHARI AND STEFANO SOATTO

ABSTRACT. We analyze the regularization properties of additive gradient noise in the training of deep networks by posing it as finding the ground state of the Hamiltonian of a spherical spin glass in an external magnetic field. We show that depending upon the magnitude of the magnetic field, the Hamiltonian changes dramatically from a highly non-convex energy landscape with exponentially many critical points to a regime with polynomially many critical points and finally, "trivializes" to exactly one minimum. This phenomenon, known as topology trivialization in the physics literature, can be leveraged to devise annealing schemes for additive noise such that the training starts in the polynomial regime but gradually morphs the energy landscape into the original one as training progresses. We demonstrate through experiments on fully-connected and convolutional neural networks that annealing schemes based on trivialization lead to accelerated training and also improve generalization error.

## 1. INTRODUCTION

Deep learning has enjoyed enormous success in recent years, in applications ranging from image recognition (Krizhevsky et al., 2012), speech recognition (Hinton et al., 2012a) and natural language processing (Bengio et al., 2003) to reinforcement learning (Mnih et al., 2015). Consequently, a number of recent papers have focused on different aspects of this problem such as the hierarchical structure (Patel et al., 2015; Salakhutdinov and Hinton, 2009; Damianou and Lawrence, 2012), robustness of representations (Chatfield et al., 2014; Nguyen et al., 2014), large scale optimization on distributed platforms (Seide et al., 2014; Dean et al., 2012) etc. This paper is in the spirit of results that analyze algorithms for training deep networks using knowledge of the landscape of the cost function, e.g., Dauphin et al. (2014); Haeffele and Vidal (2015); Janzamin et al. (2015).

We derive motivation from the theory of spin glasses. In recent years, the connections of spin glasses to combinatorial optimization, random $k$-SAT, LDPC codes etc. have become popular (Mezard and Montanari, 2009) and an understanding of the energy landscape and thermodynamics of such systems has led to a number of efficient algorithms (Krzakala et al., 2011; Braunstein et al., 2005). Spin glasses are particularly attractive models for large-scale problems such as deep networks — in fact, they are among the oldest models for neural networks (Hopfield, 1982) — because although their energy landscape is a high-dimensional, non-convex Gaussian field, we can explicitly compute various quantities of interest such as the expected number of critical points (Auffinger et al., 2013) and the structure of the Gibbs distribution (Talagrand, 2003; Mezard et al., 1987).

In this work, we consider a model for deep networks with random, sparse weights and connect its loss function to the Hamiltonian of a spherical spin glass. Existing results from statistical physics show that the Hamiltonian has a complex landscape with exponentially many critical points of all orders and stochastic gradient descent (SGD) algorithms cannot make progress (Auffinger et al., 2013; Choromanska et al., 2014). An effective way of changing this landscape is to perturb the Hamiltonian by an external magnetic field (Fyodorov, 2013). In fact, there exists a critical threshold

Computer Science, UCLA. pratikac@ucla.edu.
Computer Science, UCLA. soatto@cs.ucla.edu.

of the perturbation below which the exponentially many critical points persist and above which the landscape has only one local minimum. More importantly, this phase transition is not sharp, there exists a small band around the perturbation threshold that results in polynomially many critical points and thus gives a good trade-off between the original complex landscape and a completely degenerate loss function.

We interpret the external magnetic field as regularization of the original loss function of a deep network by additive noise and construct annealing schemes where the magnitude of the perturbation is gradually changed to transition from the polynomial regime in the beginning of training to recover the original landscape as the training progresses. We extend the idea to other types of regularizations, e.g, weight decay, dropout, multiplicative noise and propose a way to accelerate training of deep networks with these regularizations. Finally, we discuss experiments on fully-connected deep networks and convolutional neural networks that demonstrate improved training performance and generalization error.

1.1. **Related work.** Our work is closely related to Choromanska et al. (2014) which uses spin glasses to analyze the energy landscape of deep networks. Their experiments show that although there are exponentially many saddle points of the Hamiltonian, the Hessian at the optima recovered by stochastic gradient descent (SGD) has very few negative eigenvalues. Sagun et al. (2014) show that SGD fails to progress beyond a very specific energy barrier, one that can be seen in spin glass computations as the onset of high-order saddle points and local minima. In this paper, we demonstrate a way to use topology trivialization to modify this energy landscape and make it more amenable to gradient descent.

Neelakantan et al. (2015) employs annealed gradient noise to show through experiments that it helps training of large, complex recurrent networks. We provide a formal analysis of this phenomenon, indeed, gradient noise annealing is a direct application of "untrivializing" the loss function. The magnitude of noise used by Neelakantan et al. (2015) scales as $k^{-0.55}$ for the $k^{th}$ iteration, while our analysis provides an explicit formula for a variety of regularization schemes in terms of the number of neurons $n$ and layers $p$.

The random, sparse model in Sec. 2.1 is motivated from Arora et al. (2013) and Neyshabur and Panigrahy (2013). Empirically, a number of works such as Denil et al. (2013); Chen et al. (2015) have reported that as much as 95% connections in a typical deep network have almost zero weights. Such considerations have also resulted in state-of-the-art architectures, for instance, the "inception network" (Szegedy et al., 2014). The random, sparse network allows us to make the necessary approximations to get an analytical handle on its loss function, but as experiments in Sec. 5 show, the insight derived from trivialization on this model also applies to fully-connected and convolutional neural networks.

Conceptually, our approach is similar to homotopy continuation methods (Allgower and Georg, 2012; Mobahi, 2016) which start optimization with a benign energy landscape and gradually recover the original problem. Our theoretical analysis in Sec. 3 is however specialized to spin glass Hamiltonians. In our aim to trivialize the loss function and sequentially tighten this relaxation, our work also relates to convex relaxation strategies used in machine learning, e.g., maximum-a-posteriori estimation (Jaakkola et al., 2010) and sparse principal component analysis (Wang et al., 2014).

## 2. Preliminaries

2.1. **A model for deep networks.** Let us consider a deep network with $p$ hidden layers and $n$ neurons on each layer. We denote the observed layer at the bottom by $X \in \{0,1\}^n$ which is generated

by $h \in \{0,1\}^n$ at the top. More precisely, $X$ can be written as

$$X = g\left(W^{1^\top}\, g\left(W^{2^\top} \dots\, g\left(W^{p^\top}\, h\right)\dots\right)\right); \tag{1}$$

where $W^k \in [-N,\, N]^{n\times n}$ for $k \le p$ and some constant $N$ are the connectivity matrices and $g(x) = \mathbf{1}_{\{x\ge 0\}}$ are threshold nonlinearities. $\mathbf{1}_A$ denotes the indicator function of the set $A$. The output of the intermediate layers is given by $h^{k-1} = g\left(W^{k^\top}\, g\left(W^{k+1^\top} \dots\, g\left(W^{p^\top}\, h\right)\dots\right)\right)$ where $h^{k-1} \in \{0,1\}^n$. With this notation, we have $X = h^0$ and $h = h^p$. To make this model analytically tractable, we make the following assumptions:

(1) $h$ has at most $\rho$ non-zero entries,
(2) For $d = n^{1/p}$, every entry $W^k_{ij}$ is an iid zero-mean random variable with $\mathrm{P}\left(W^k_{ij} > 0\right) = \mathrm{P}\left(W^k_{ij} < 0\right) = d/(2n)$ and zero otherwise.
(3) The cardinality of the support of $W^k_{ij}$ is $m = \Theta(n)$ and the weight distribution is not concentrated near zero (cf. Appendix A) and that $\rho\left(\frac{d}{2}\right)^p$ is a constant.

If every entry of $h_i$ is an indicator of some class, the first assumption implies that at most $\rho$ classes are present in every data sample $X_i$. The Chernoff bound then implies that with high probability, each $X_i$ has $\rho\left(\frac{d}{2}\right)^p$ fraction non-zero entries, which we assume to be constant. The third assumption on the support of the weight distribution ensures that the weights of the deep network are well-conditioned.

The above model enjoys a number of useful properties by virtue of it being random and sparsely connected. For high sparsity ($d < n^{1/5}$ or equivalently, $p > 5$), the network exhibits "weight-tying", i.e., it acts as a denoising auto-encoder using the same weights

$$h = g\left(W^p g\left(W^{p-1} \dots g\left(W^1 \widetilde{X} - \frac{d}{3}\mathbb{1}_n\right) - \frac{d}{3}\mathbb{1}_n\right)\dots\right);$$

where $\mathbb{1}_n$ denotes a vector of ones of length $n$ and $\widetilde{X} = X + \delta$ for some small noise $\delta$. In other words, the feature vector $h$ can be obtained from $X$ by running the same network in reverse with a different thresholding function $g'(x) = \mathbf{1}_{\left\{x \ge \frac{d}{3}\right\}}$ Arora et al. (2013). We will use this property to compute the distribution of the loss function of this network.

**Classification model:** In a typical classification scenario, one uses a support vector machine (SVM) or a soft-max layer on the feature vector $h$ to obtain an output $Y \in \{0,1\}$. We model this as

$$Y = g'\left(W^{p+1} g'\left(W^p \dots g'\left(W^1 X\right)\dots\right)\right); \tag{2}$$

where $g'(x) = \mathbf{1}_{\left\{x \ge \frac{d}{3}\right\}}$ is a thresholding function and $W^{p+1} \in [-N,\, N]^n$ is a $d$-sparse row vector. The data sample $X \in \{0,1\}^n$ is denoted as the input while the output is $Y \in \{0,1\}$. We assume that the true label $Y^t$ is independent of $X$. We will henceforth consider (2) as the canonical model of a deep neural network.

2.2. **Deep networks as spin glasses.** A $p$-spin glass for an integer $p \ge 1$ with $n$ spins is given by an energy function

$$-H_{n,p}(\sigma) = \frac{1}{n^{(p-1)/2}} \sum_{i_1,i_2,\dots,i_p=1}^{n} J_{i_1,\dots,i_p}\sigma_{i_1}\dots\sigma_{i_p}; \tag{3}$$

where $\sigma = (\sigma_1,\dots,\sigma_n)$ is a configuration of the spins and $J_{i_1,\dots,i_p}$ are iid, zero-mean, standard Gaussian random variables. The energy function is usually called the Hamiltonian while the coefficients $J_{i_1,\dots,i_p}$ are known as the disorder. Note that $H_{n,p}(\sigma)$ is zero-mean and Gaussian for every $\sigma$. The constant $n^{-(p-1)/2}$ is present to ensure that the Hamiltonian is extensive, i.e., it scales with the number

of spins $n$. We now use the following lemma to connect the model for deep networks in Sec. 2.1 with the spin glass Hamiltonian.

**Lemma 1.** *If the true label $Y^t \sim Ber(q)$ for some $q < 1$ is independent of data, the loss function defined as $H_{n,p} = (1-2q)^{-1}\left(1 - |\widehat{Y}, Y^t|\right)$ is given by*

$$-H_{n,p}(\sigma) = \frac{J}{n^{(p-1)/2}} \sum_{i_1,i_2,\ldots,i_p=1}^{n} J_{i_1,\ldots,i_p}\, \sigma_{i_1} \ldots \sigma_{i_p}; \tag{4}$$

*where $J_{i_1,\ldots,i_p}$ is a zero-mean, standard Gaussian random variable, $J \in \mathbb{R}$ is a constant and $\sigma \in \mathbb{R}^n$.*

*Proof.* (sketch, cf. Appendix A) We write the output $Y$ as a function of the active paths from input neurons $X_i$s and show that the correlation introduced in these paths due to the inputs is zero-mean with variance $\Theta(n^{-2/5})$. The probability of a path being active in the sparse network is $\Theta(n^{-(p-1)/2})$. We approximate these two as the Gaussian disorder $J_{i_1,\ldots,i_p}$. Moreover, since the weights are supported on a finite set, we can write the weight of each path in terms of the unique weights in the network denoted by $\sigma \in [-N,N]^n$. ∎

We emphasize that in the sequel, $\sigma$ is the quantity that characterizes the weights of our deep network. We additionally assume a spherical constraint on the spins, i.e., $\frac{1}{n}\sum_i \sigma_i^2 = 1$ or equivalently, $\sigma \in S^{n-1}(\sqrt{n}) \subset \mathbb{R}^n$.

## 3. Perturbations of the Hamiltonian

This section discusses the effect of external perturbations on the spin glass Hamiltonian. If the perturbations are large, the landscape is dominated by the perturbation term resulting in significantly fewer critical points, while if the magnitude is small, the Hamiltonian retains its original landscape with exponentially many critical points.

Let the total number of critical points in the set $(-\infty, u]$ be $\mathrm{crt}(u) = |\{\sigma :\ \nabla H(\sigma) = 0,\ H(\sigma) \leq n\, u\}|$. With minor abuse of notation, denote $\mathrm{crt}(\infty)$ by $\mathrm{crt}(H)$. The index of a critical point is the number of negative eigenvalues of the Hessian of $H_{n,p}(\sigma)$ and we denote critical points of index $k$ by $\mathrm{crt}_k(u)$. In this notation, local minima are therefore given by $\mathrm{crt}_0(u)$. As Auffinger et al. (2013) show, we can explicitly compute the limiting value of $\mathrm{crt}(u)$ and $\mathrm{crt}_k(u)$. For any $k \geq 0$ and $p \geq 2$, we have

$$\lim_{n\to\infty} n^{-1} \log \mathrm{E}\, \mathrm{crt}_k(u) = \Theta_k(u), \tag{5}$$

$$\lim_{n\to\infty} n^{-1} \log \mathrm{E}\, \mathrm{crt}(u) = \Theta(u); \tag{6}$$

where $\Theta_k(u)$ and $\Theta(u)$ are functions known as the complexity of a spin glass (cf. Fig. 1). The notation "E" denotes expectation and we use $\mathrm{E}\, X$ as a shorthand for $\mathrm{E}(X)$.

### 3.1. Scaling of critical points.

In the non-asymptotic regime, Thm. 2 characterizes the number of critical points of the Hamiltonian in terms of the density of eigenvalues of the Gaussian orthogonal ensemble (GOE); see Appendix B for a brief introduction.

**Theorem 2** (Fyodorov (2013))**.** *For a zero-mean Hamiltonian $H(\sigma)$ with covariance $E\left(H(\sigma)H(\sigma')\right) = F(\sigma^\top \sigma')$ that is invariant under the transformations $\sigma \mapsto O\sigma$, $\sigma' \mapsto O\sigma'$ for $\sigma \in S^{n-1}(\sqrt{n})$ and any orthogonal matrix $O \in \mathbb{R}^{n\times n}$, the expected number of critical points of $H$ is*

$$E\, crt(H) = 4n \left(\frac{1+B}{1-B}\right)^{n/2} \sqrt{1-B}\; G(B);$$

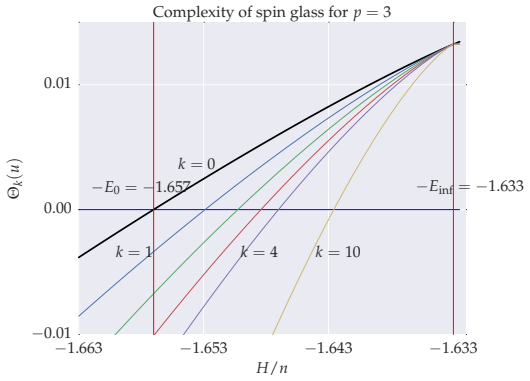


FIGURE 1. Roughly, the number of critical points in the set $(-\infty, u]$ scales as $e^{n\Theta_k(u)}$. The black curve denotes the complexity of local minima, $\Theta_0(u)$. Below $H/n = -E_0$, there are no local minima in the logarithmic scaling as $n \to \infty$. It can be shown that $n^{-1}\inf H(\sigma)$ converges to $-E_0$ in probability (Auffinger et al., 2013); we denote this as the energy barrier at $-E_0$. Similarly, there exists another barrier $-E_{\text{inf}}$ of the normalized Hamiltonian beyond which the index $k$ of critical points diverges with $n$.

*where*

$$B = \frac{F''(n) - F'(n)/n}{F''(n) + F'(n)/n}, \qquad G(B) = \int_0^\infty e^{-nBt^2/2}\, \rho_n(t)\, dt,$$

*and $\rho_n(t)$ is the density of the eigenvalues of GOE.*

We now apply the above theorem to a perturbed $p$-spin glass Hamiltonian in (4) given by

$$-\widetilde{H}(\sigma) = \sum_{i_1, i_2, \ldots, i_p = 1}^{n} J_{i_1,\ldots,i_p}\, \sigma_{i_1} \ldots \sigma_{i_p} + \sum_{i=1}^{n} h_i\, \sigma_i; \tag{7}$$

where $J_{i_1,\ldots,i_p}$ is a zero-mean Gaussian random coupling with variance $J^2\, n^{-p+1}$ and $\sum_i h_i \sigma_i$ is the external perturbation. We let $h_i$ be iid zero-mean Gaussian random variables with variance $\mathrm{E}\,(h_i^2) = \nu^2$. Thm. 3 shows that the parameter $\nu$ results in three distinct regimes for the expected number of critical points of the Hamiltonian in (7) with

$$B = \frac{J^2 p(p-2) - \nu^2}{J^2 p^2 + \nu^2} \in \Big(-1, \frac{p-2}{p}\Big]. \tag{8}$$

**Theorem 3** (Fyodorov (2013), Sec. 2.3)**.** *For $p > 2$ and large $n$, the expected number of critical points of the perturbed Hamiltonian is given by*

$$E\, crt(\widetilde{H}) = \begin{cases} 2 & \text{if } B = -\Omega(n^{-1}), \\ \frac{2n}{\sqrt{\pi}}\tau^{-3/2} & \text{if } B = -\frac{\tau}{n}, \\ 4n^{1/2}\sqrt{\frac{1+B}{\pi B}}\exp\left(\frac{n}{2}\log\frac{1+B}{1-B}\right) & \text{if } B > 0. \end{cases}$$

*where $B = 0$ corresponds to the critical perturbation magnitude $\nu_c = J\sqrt{p(p-2)}$ and $\tau \gg 1$ is a constant.*

The first case above shows total trivialization because any smooth function on the sphere $S^{n-1}(\sqrt{n})$ has at least one minimum and at least one maximum, thus there is exactly one minimum if the perturbations are larger than $J\sqrt{p(p-2)}$. For $\nu < \nu_c$, we are back to the exponential regime given by Thm. 2. There is also an edge-scaling regime around the value $B = 0$ where if $B = -\mathcal{O}(1/n)$ the

spin glass has $\mathcal{O}(n)$ critical points. We would like to pick the perturbation magnitude to be in this regime to not only avoid exponentially many critical points, but also avoid total trivialization.

**Corollary 4.** *For the model in Sec. 2.1 for $p \gg 2$ and large $n$, we need $\nu = Jp(1+2\tau/n)^{1/2}$ to be in the polynomial regime given by Thm 3.*

The above corollary follows directly from (8) and the fact that Thm. 3 requires $B = -\tau/n$ for the polynomial regime. The following lemma describes the picture near the perturbation threshold and makes the role of $\tau$ more precise.

**Lemma 5** (Fyodorov (2013), Eqn. 67)**.** *For $B = -\tau/n$ and $\tau \in \mathbb{R}$, the expected number of critical points is given by*

$$\lim_{n\to\infty} \frac{E\, crt(\widetilde{H})}{2n} = \begin{cases} \frac{1}{\sqrt{\pi}}\, \tau^{-3/2} & \text{for } \tau \gg 1, \\ \frac{1}{\sqrt{\pi|\tau|}}\, 2e^{|\tau|} & \text{as } \tau \to -\infty. \end{cases} \tag{9}$$

*Note that the $\tau \gg 1$ case in (9) matches the polynomial regime in (3) while the $\tau \to -\infty$ case matches the $B \ll 1$ limit of the exponential regime in (3).*

3.2. **Quality of perturbed local minima.** The expected number of local minima of the Hamiltonian also undergoes a similar phase transition upon adding an external perturbation with the same critical threshold, viz., $\nu_c = J\sqrt{p(p-2)}$. If $\nu < \nu_c$, the energy landscape is relatively unaffected and we have exponentially many local minima as predicted by (5) and exactly one minimum in expectation otherwise. Moreover, there exist polynomially many local minima in the band $|B| = \mathcal{O}(n^{-1/3})$. Appendix C discusses the corresponding versions of Thm. 3 and Lem. 5.

Thus far, we have only computed the number of critical points of the perturbed Hamiltonian, in particular, we do not have bounds on the value of the Hamiltonian at these perturbed critical points. The following lemma shows that small perturbations, viz., $\|h\| = \mathcal{O}(n^{1/2})$, do not change the locations of minima much, and the value of the normalized Hamiltonian $H(\sigma)/n$ due to such perturbations changes by at most $2\nu$. Appendix D has a more precise formulation.

**Lemma 6.** *For $\nu = \mathcal{O}(1)$ and large enough $n$, local minima are perturbed by at most $\mathcal{O}(n^{-1/3})$ while the normalized Hamiltonian at these minima changes by at most $2\nu$.*

3.3. **Annealing scheme.** We leverage upon this analysis in the following way. As Lem. 5 shows, the parameter $\tau$ can smoothly interpolate between the polynomial regime ($\tau \gg 1$) and the exponential regime ($\tau \to -\infty$). We can now construct an annealing procedure which reduces the perturbation by annealing $\tau$ and gradually changes the Hamiltonian from the polynomial regime to the original energy landscape as training progresses. Since the perturbation goes to zero by the end of training, the quality of local minima after such an annealing is unaffected. Moreover, by Lem. 6, even during annealing, the local minima are only perturbed by a small amount.

If $N$ is the total number of training epochs, and $k$ is the current epoch, we set $\tau(k) = -\frac{nk}{2N}$ (cf. Cor. 4). Note that this gives $\nu = Jp$ for $k = 0$ and $\nu = 0$ for $k = N$. Also, for $k \ll N$, the perturbation $\nu$ scales as $(1+N^{-1}k)^{-1/2}$ which matches the empirical results in Neelakantan et al. (2015).

## 4. Regularization in deep networks

We now interpret different types of regularization in deep networks as perturbations of the original loss function (4) and compute the scaling of additive noise in terms of the number of neurons $n$ and the number of layers $p$ for these regularizations, to realize the polynomial regime of the energy landscape. Conversely, one can invert the formulas in this section to compute the appropriate regularization coefficients for a given level of additive noise, for instance, in a hyper-parameter tuning setting (Bergstra and Bengio, 2012).

4.1. **Additive noise.** The weight update equation in the presence of additive noise (Jim et al., 1996) can be written as

$$\sigma \leftarrow \sigma - \eta \, \nabla H_{n,p}(\sigma) - \Delta;$$

where $\Delta \in \mathbb{R}^n$ is a zero-mean Gaussian random vector with variance $v^2 I_{n\times n}$, $\nabla H_{n,p}(\sigma)$ is the back-propagated gradient and $\eta$ is the learning rate. Additive noise is typically used to escape local minima and thus it is important to pick the right magnitude of var $\Delta$ to avoid large jumps in the weights that are uncorrelated with the gradient. The perturbed Hamiltonian that gives the weight update equation above, $-\widetilde{H}_{\text{add}}(\sigma) = -H_{n,p}(\sigma) + \sum_{i=1}^{n} \Delta_i \, \sigma_i$, is the same as (7) and for a constant $\tau$, we have from Cor. 4 that

$$v = J \; p \left(1 + \frac{2\tau}{n}\right)^{1/2}. \tag{10}$$

4.2. **Weight decay.** Weight decay (Moody et al., 1995) implies an $\ell_2$ regularization on the weights or equivalently, a Gaussian prior on them. For a constant $\alpha > 0$, we write the weight updates as

$$\sigma \leftarrow \sigma - \eta \, \nabla H_{n,p}(\sigma) - pn^{-1/2} \alpha \, \sigma;$$

the factor $pn^{-1/2}$ is merely for convenience. The decay term above can be obtained using an additive term of the form $\frac{p\alpha}{2\sqrt{n}} \sum_{i=1}^{n} \sigma_i^2$ in the Hamiltonian. Weight decay by itself will not lead to any phase transitions in the number of critical points. However, if we add a perturbation term in addition to $\ell_2$ regularization, we can expect similar behavior as Sec. 3. Consider the zero-mean perturbation

$$-\widetilde{H}_{\text{decay}}(\sigma) = -H_{n,p}(\sigma) + \frac{p\alpha}{2\sqrt{n}} \sum_{i=1}^{n} \sigma_i^2 + \sum_i h_i \, \sigma_i \tag{11}$$

with var $h_i = v^2$ which creates an additional additive term of $h$ in the weight update rule above. $\widetilde{H}_{\text{decay}}$ is now amenable to the analysis of Thm. 3 with the covariance given by

$$\begin{aligned} \frac{J^2}{n^{p-1}} \sum_{i_1, i_2, \ldots, i_p} \prod_{k=1}^{p} (\sigma_{i_k} \sigma'_{i_k}) + \frac{p^2\alpha^2}{4} \sum_{i,j} \left(\sigma_i \sigma'_j\right)^2 + v^2 \sum_i \sigma_i \sigma'_i, \\ \approx n \left( J^2 u^p + \frac{p^2\alpha^2}{4} u^2 + v^2 u \right), \\ \triangleq n f\left(\frac{\sigma^T \, \sigma'}{n}\right); \end{aligned}$$

where $u = \left(\sigma^T \sigma'\right)/n$. The approximation in the second step is obtained by assuming that $\mathrm{E}(\sigma_i \sigma'_j) \approx 0$ for all $i \neq j$ which roughly says, averaged over the disorder $J$, two deep networks given by $\sigma, \sigma'$ have uncorrelated $i^{\text{th}}$ and $j^{\text{th}}$ neurons on each layer. We now use the above expression to calculate the value of $B$ in (8) to get $B = \frac{J^2 p(p-2) - v^2}{J^2 p^2 + p^2 \alpha^2 + v^2}$. Note that the value of $v_c$ which corresponds to $B = 0$ is unchanged from Thm. 3, but to stay in the polynomial regime we need

$$v = Jp \left(1 + \tau \, \frac{\alpha^2}{n} + \frac{2\tau}{n}\right)^{1/2}. \tag{12}$$

Contrast this with Cor. 4 or additive noise in Sec. 4.1; the strength of the perturbation needs to be larger to account for the curvature of the $\ell_2$ regularization term.

The annealing procedure in Sec. 3.3 becomes

$$\tau(k) = -\left(\frac{n}{\alpha^2 + 2}\right) \frac{k}{N}. \tag{13}$$

Let us emphasize that we fix the perturbation $h$ at the beginning of the training and do not resample it. At each epoch, we simply rescale its magnitude using (13). The motivation in literature behind resampling additive gradient noise at each iteration is to enable gradient descent to jump out of local minima. In this work, we instead rely on the weights aligning themselves in the direction of the magnetic field in the beginning of training. As the magnitude of this external field reduces, the properties of stochastic gradient descent dominate the training performance (cf. Figs. 4b and 4c).

4.3. **Multiplicative noise.** If instead of a constant $\alpha$, we use a random matrix $\Delta \in \mathbb{R}^{n\times n}$ (usually diagonal) with iid entries $\mathrm{E}\,\Delta_{ii} = 0$ and $\mathrm{E}\,(\Delta_{ii}^2) = \alpha^2$ to multiply the weights, we get multiplicative noise (Nalisnick et al., 2015) with the update rule

$$\sigma \leftarrow \sigma - \eta\,\nabla H_{n,p}(\sigma) - pn^{-1/2}\,\Delta\,\sigma.$$

Similar to weight decay, we can add a perturbation term of $\frac{p}{2\sqrt{n}}\,\sum_{i=1}^n\,\Delta_{ii}\,\sigma_i^2$ to the Hamiltonian. The analysis now follows that of Sec. 4.2 and we get the same expression as (12) although since entries of $\Delta$ are iid, we need not make the approximation that $\mathrm{E}\,(\sigma_i\sigma_j') \approx 0$ for all $i \neq j$.

Appendix E derives similar expressions for additive noise for DropOut (Hinton et al., 2012b) and Drop-Connect (Wan et al., 2013) regularizations.

## 5. Experiments

We now discuss simulation experiments on the energy landscape of perturbed spin glass Hamiltonians. We also demonstrate how the annealing scheme obtained in Sec. 3.3 can be used to improve training performance and generalization error on fully-connected and convolutional neural networks.

5.1. **Spin glasses.** Consider a 3-spin glass Hamiltonian, which in our notation is a fully-connected network with three hidden layers, threshold nonlinearities and a fourth loss layer, given by

$$-\widetilde{H}_{n,3} = \frac{1}{n}\sum_{i,j,k=1}^{n} J_{i,j,k}\,\sigma_i\sigma_j\sigma_k + \sum_i h_i\,\sigma_i; \tag{14}$$

with $h_i \in N(0,\nu^2)$, $\sigma \in S^{n-1}(\sqrt{n})$ and $J_{i,j,k} \sim N(0,1)$. Given a disorder $J$ we find the local minima of the above Hamiltonian for values of $h$ in the three regimes considered in this paper (cf. Cor. 4): (i) $\nu = 1/n$ for exponentially many critical points, (ii) $\nu = (1+1/n)^{1/2}$ to be in the polynomial regime and, (iii) $\nu = p$ to be in the totally trivialized regime with a single minimum. For an initial configuration $\sigma^1$ that is uniform on $S^{n-1}(\sqrt{n})$, we perform gradient descent with a fixed step size $\eta = 0.1$ and stop when the gradient $\|\nabla\widetilde{H}\|_2 \leq 10^{-4}$. Note that at each step, we need to project the gradient on the tangent plane of $S^{n-1}(\sqrt{n})$ and re-project $\sigma^{k+1} \leftarrow \sigma^k - \eta\nabla\widetilde{H}(\sigma^k)$ back on to the sphere.

In order to visualize the high-dimensional energy landscape, we perform a t-distributed stochastic neighbor embedding (t-SNE) (Van der Maaten and Hinton, 2008) of the local minima obtained by starting gradient descent from 20,000 uniformly sampled points on $S^{n-1}(\sqrt{n})$ for $n = 100$. t-SNE is a nonlinear dimensionality reduction technique that is capable of capturing the local structure of the high-dimensional data while also revealing global structure such as the presence of clusters at several scales. We use the Euclidean metric on $\mathbb{R}^n$ as a measure of similarity between points for t-SNE. In particular, we do not use the value of the Hamiltonian for t-SNE. Fig. 2 discusses the results of this simulation.

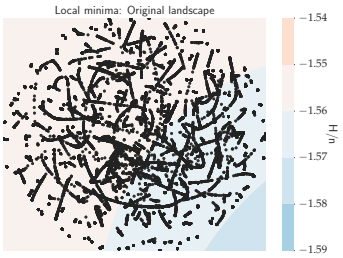


(A) Exponential regime: In addition to local minima surrounded by high energy barriers (isolated dots), we also see "canyons" with long narrow passages. This suggests that there are large regions where the gradient is very small in all but a few directions; these are thus saddle points.

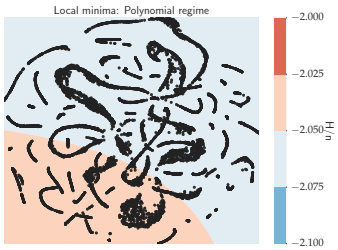


(B) Polynomial regime: As Thm. 3 predicts, the number of critical points seen as isolated clusters in the t-SNE projection is significantly reduced as compared to Fig. 2a. Surprisingly, the landscape is full of saddle points in this regime.

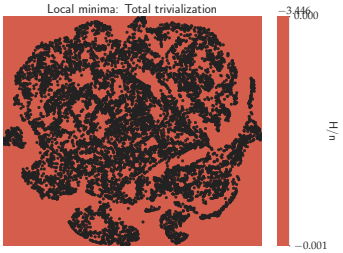


(C) Total trivialization: All executions of gradient descent converge at the same local minimum with the almost same value of normalized Hamiltonian (−3.4463). The average distance between two points in the original high-dimensional data here is 0.02 as compared to 1.16 for Fig. 2a. t-SNE thus projects all points into a single cluster which is an indication of a unique local minimum.

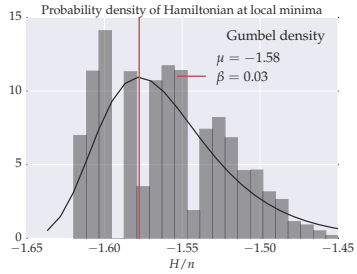


(D) For $n = 100$, roughly, $H(\sigma)/n$ at local minima follows an extreme-value distribution (negative Gumbel) given by $p(x) = \beta^{-1}\exp(z - e^z)$ where $z = -(x-\mu)/\beta$ with $\mu \approx -1.62$ and $\beta \approx 0.006$ (Subag and Zeitouni, 2015). We obtain the parameters $\mu = -1.58$ and $\beta = 0.03$, which suggests that with a small magnetic field ($\nu = 1/n$), we recover the distribution of the local minima.

FIGURE 2. Two-dimensional t-SNE (Van der Maaten and Hinton, 2008) of 20,000 local minima on the sphere $S^{n-1}(\sqrt{n})$ along with the values of the Hamiltonian obtained by gradient descent on a 3-spin glass with $n = 100$ in the three regimes, exponential (Fig. 2a), polynomial (Fig. 2b) and trivial (Fig. 2c). We remark that Figs. 2a and 2d use the exact same data, but the background in Fig. 2a is colored with the kernel density estimate of the Hamiltonian which misses the tails of the distribution (cf. Fig. 2d) due to less data there. Also, the normalized Hamiltonian in Figs. 2b and Fig. 2c has an large additive term, viz., $\frac{1}{n}\left|h^\top \sigma\right| \approx \mathscr{O}(\nu)$.

5.2. **Fully-connected deep networks.** We now use the annealing procedure described in Sec. 3.3 along with weight decay to train fully-connected deep networks (FcNNs) for image classification on two datasets, viz., the MNIST dataset (LeCun et al., 1998) with $32 \times 32$ gray-scale images of digits, and the CIFAR-10 dataset (Krizhevsky, 2009) with $32 \times 32$ RGB images. We whiten both datasets by subtracting the mean and dividing by the per-pixel standard deviation, no data augmentation is performed. Both networks have $p = 20$ layers with relu nonlinearities and use soft-max over 10 classes as a classifier. The mnistfc network has 64 hidden units on each layer while the cifarfc network has 128 hidden units on each layer. We use a momentum-based optimization algorithm

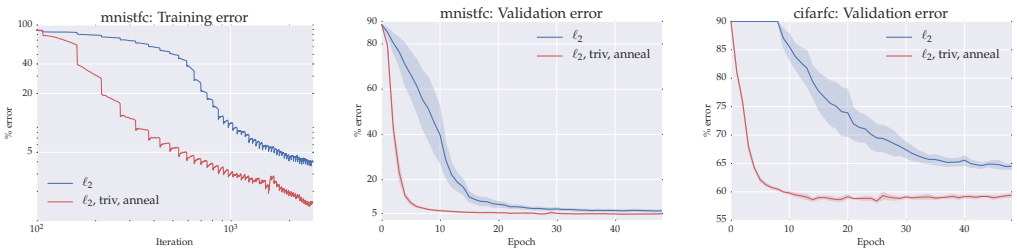


(A) (B) (C)

FIGURE 3. Figs. 3b and 3c show the validation error averaged over 10 runs of 50 epochs and a batch size of 1024 for mnistfc and cifarfc. Fig. 3a shows the training error for mnistfc (cifarfc is similar). Blue denotes an $\ell_2$ regularization with $\alpha = 10^{-8}$ without additive noise while red denotes the same $\ell_2$ regularization and additive noise annealing with $J = 10^{-5}$. Trivialization accelerates training performance dramatically, the networks are very close to their final validation error by the $5^{th}$ epoch, 4.9% for mnistfc and 59.8% for cifarfc. These networks lack convolutions and hence we do not expect state-of-the-art validation errors. Annealing also helps initialization of weights in deep FcNNs, for instance, the validation error in Fig. 3c remains at 90% (same as guessing uniformly) for almost 9 epochs.

ADAM (Kingma and Ba, 2014) after tuning the initial learning rate to $\eta = 10^{-3}$ that decays by a factor of 0.95 after every epoch and saturates at $10^{-5}$. Fig. 3 discusses the results on fully-connected deep networks.

We estimate the number of neurons on each layer as $n = \sqrt{\#\text{weights}/p}$ for the expressions of gradient noise (12) and annealing (13). This estimate is accurate if all layers are fully-connected and of the same size, but we use the same expression for experiments on CNNs with good results. Also, note that the model in Sec. 2.1 has $p$ hidden layers in addition to the loss layer, so we will consider each linear (or convolutional) layer with nonlinearities, pooling etc. as one layer. In practice, we treat the constant $J$ as a tunable parameter and set its value using cross-validation.

5.3. **Deep convolutional neural networks.** We next train the following CNN on the CIFAR-10 dataset:

$$\text{cifarconv}: \underbrace{\text{conv}_{5\times5\times8}}_{\times\,8} \to \underbrace{\text{linear}_{128}}_{\times\,8} \to \text{linear}_{10} \to \text{soft-max};$$

with a max-pool$_{3\times3}$ layer after each convolutional layer and relu nonlinearities after every layer. The notation conv$_{5\times5\times8}$ denotes a convolutional layer with a $5\times5$ kernel, 8 output features and the appropriate number of input features. We use zero-padding on convolutional and max-pooling layers to preserve the size of the input through the convolutional part of the network. The number of layers for cifarconv is $p = 16$. Fig. 4 describes the results of this section.

5.4. **Network in network architecture.** The above experiments show that annealing additive noise accelerates training and improves generalization error on "narrow and deep" networks. We now discuss the results on the CIFAR-10 dataset using a network with a large number of hidden units but a smaller depth denoted by

$$\text{cifarnin}: \underbrace{\underbrace{\text{conv}_{5\times5\times192}}_{\times\,3} \to \text{max-pool}_{3\times3,\text{dropout}}}_{\times 2} \to \text{conv}_{3\times3\times192} \to$$
$$\ldots \text{conv}_{1\times1\times192} \to \text{conv}_{1\times1\times10} \to \text{mean-pool}_{8\times8} \to \text{soft-max}.$$

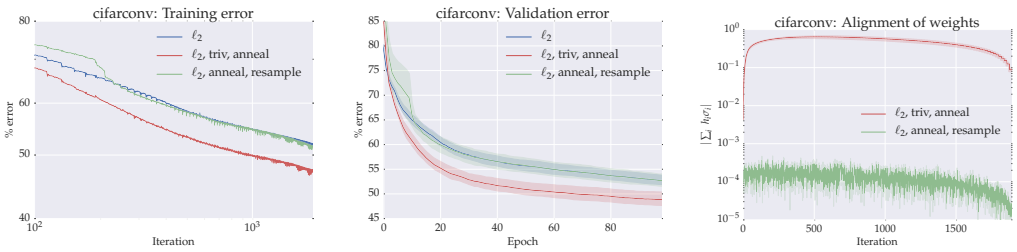


(A) (B) (C)

FIGURE 4. Figs. 4a and 4b show the training and validation error of the cifarconv network averaged over 5 runs of 100 epochs with a batch size of 256. Blue denotes weight decay with $\alpha = 10^{-8}$ and red denotes an additional annealed additive noise term with $J = 10^{-5}$. The green curve employs the same weight decay and annealing but resamples additive noise at each iteration. Trivialization accelerates the training in the beginning as the weights align themselves with the external magnetic field (cf. Fig. 4c). As the perturbations are annealed, $|\sum_i h_i \sigma_i|$ decreases towards zero as well. Resampling additive noise at each iteration performs worse than trivialization because the weights are forced to align with the randomly sampled magnetic field at each iteration.

Every convolutional layer in the "network-in-network" architecture (Lin et al., 2013) above is followed by batch-normalization (Ioffe and Szegedy, 2015) and relu nonlinearities and classification is performed by global average pooling (Lin et al., 2013). These regularization techniques are designed to avoid over-fitting of fully-connected layers and this architecture has resulted in state-of-the-art performance on a number of datasets. We set the number of layers $p = 8$, one for each convolutional layer before the final global average pooling classifier. We remark that the cifarnin network is very far from the model considered in Sec. 2.1 but as Fig. 5 shows, we can still reap the benefits of additive noise annealing while training this network.

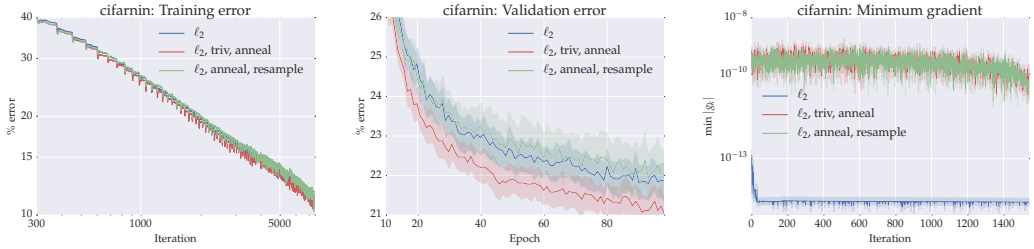


(A) (B) (C)

FIGURE 5. Figs. 5a and 5b show training and validation performance of the cifarnin network averaged over 5 runs of 100 epochs and a batch size of 32. As before, blue, red and green curves share $\ell_2$ regularization with $\alpha = 10^{-8}$, the red curve employs annealed additive noise with $J = 10^{-6}$ while the green curve has the same magnitude annealing but resamples noise at each iteration. The heavy regularization introduced by the network-in-network architecture results in smaller gains due to trivialization; nevertheless the blue curve has an average testing error of 21.89% as opposed to 21.02% average error on the red curve with trivialization and 22.05% average error on the green curve. Additive noise ensures that the gradient is bounded away from zero, we check this by plotting the smallest absolute value in the back-propagated gradient in Fig. 5c. Gaussian additive noise, with or without annealing, has a much larger smallest gradient than the blue curve, but similar to Fig. 4, resampling performs worse than annealing a perturbation in a fixed direction.

## 6. DISCUSSION

We have introduced a novel topology trivialization procedure for interpreting additive noise regularization in deep learning. We can analytically compute a critical perturbation magnitude below which the loss function retains its complexity with exponentially many critical points and above which, it becomes trivial. Around this threshold, the energy landscape possesses only polynomially many critical points. This enabled us to devise an annealing procedure for additive gradient noise that gradually modifies the loss function from the polynomial regime in the beginning of training towards the original energy landscape. This annealing scheme was derived on a theoretical spin glass model, but is also applicable to fully-connected deep networks and convolutional neural networks used in practice. We found that topology trivialization results in accelerated training and also improves generalization error.

In practice, factors such as batch size and Hessian estimation parameters in SGD are responsible for gradient noise. While a systematic analysis of their effect requires further investigation, our experiments show that simply adding Gaussian noise of the right magnitude — the one given by spin glass computations — to the back-propagated gradient is extremely effective in training. Very deep networks with complex architectures which often suffer from poor initialization and vanishing gradients benefit most from this. Our analysis suggests that annealing the magnitude of perturbations, which is seen as reducing the strength of an external magnetic field, works much better than resampling the noise at each iteration, which is common in the literature.

## REFERENCES

Allgower, E. L. and Georg, K. (2012). *Numerical continuation methods: An introduction*, volume 13. Springer Science & Business Media.

Arora, S., Bhaskara, A., Ge, R., and Ma, T. (2013). Provable bounds for learning some deep representations. *arXiv:1310.6343*.

Auffinger, A., Arous, G. B., and Černỳ, J. (2013). Random matrices and complexity of spin glasses. *Communications on Pure and Applied Mathematics*, 66(2):165–201.

Bengio, Y., Ducharme, R., Vincent, P., and Janvin, C. (2003). A neural probabilistic language model. *The Journal of Machine Learning Research*, 3:1137–1155.

Bergstra, J. and Bengio, Y. (2012). Random search for hyper-parameter optimization. *The Journal of Machine Learning Research*, 13(1):281–305.

Braunstein, A., Mézard, M., and Zecchina, R. (2005). Survey propagation: An algorithm for satisfiability. *Random Structures & Algorithms*, 27(2):201–226.

Chatfield, K., Simonyan, K., Vedaldi, A., and Zisserman, A. (2014). Return of the devil in the details: Delving deep into convolutional nets. *arXiv:1405.3531*.

Chen, W., Wilson, J. T., Tyree, S., Weinberger, K. Q., and Chen, Y. (2015). Compressing neural networks with the hashing trick. *arXiv:1504.04788*.

Choromanska, A., Henaff, M., Mathieu, M., Arous, G. B., and LeCun, Y. (2014). The loss surface of multilayer networks. *arXiv:1412.0233*.

Damianou, A. C. and Lawrence, N. D. (2012). Deep gaussian processes. *arXiv:1211.0358*.

Dauphin, Y. N., Pascanu, R., Gulcehre, C., Cho, K., Ganguli, S., and Bengio, Y. (2014). Identifying and attacking the saddle point problem in high-dimensional non-convex optimization. In *Advances in Neural Information Processing Systems*, pages 2933–2941.

Dean, J., Corrado, G., Monga, R., Chen, K., Devin, M., Mao, M., Senior, A., Tucker, P., Yang, K., Le, Q. V., et al. (2012). Large scale distributed deep networks. In *Advances in Neural Information Processing Systems*, pages 1223–1231.

Denil, M., Shakibi, B., Dinh, L., de Freitas, N., et al. (2013). Predicting parameters in deep learning. In *Advances in Neural Information Processing Systems*, pages 2148–2156.

Fyodorov, Y. V. (2013). High-dimensional random fields and random matrix theory. *arXiv:1307.2379*.

Haeffele, B. D. and Vidal, R. (2015). Global optimality in tensor factorization, deep learning, and beyond. *arXiv:1506.07540*.

Hinton, G., Deng, L., Yu, D., Dahl, G. E., Mohamed, A.-r., Jaitly, N., Senior, A., Vanhoucke, V., Nguyen, P., Sainath, T. N., et al. (2012a). Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups. *Signal Processing Magazine, IEEE*, 29(6):82–97.

Hinton, G. E., Srivastava, N., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R. R. (2012b). Improving neural networks by preventing co-adaptation of feature detectors. *arXiv:1207.0580*.

Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the national academy of sciences*, 79(8):2554–2558.

Ioffe, S. and Szegedy, C. (2015). Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv:1502.03167*.

Jaakkola, T., Sontag, D., Globerson, A., and Meila, M. (2010). Learning Bayesian network structure using LP relaxations. In *International Conference on Artificial Intelligence and Statistics*, pages 358–365.

Janzamin, M., Sedghi, H., and Anandkumar, A. (2015). Beating the Perils of Non-Convexity: Guaranteed Training of Neural Networks using Tensor Methods. *arXiv:1506.08473*.

Jim, K.-C., Giles, C. L., and Horne, B. G. (1996). An analysis of noise in recurrent neural networks: convergence and generalization. *IEEE Transactions on Neural Networks*, 7(6):1424–1438.

Kingma, D. and Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv:1412.6980*.

Krizhevsky, A. (2009). Learning multiple layers of features from tiny images. Master's thesis, Computer Science, University of Toronto.

Krizhevsky, A., Sutskever, I., and Hinton, G. E. (2012). Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105.

Krzakala, F., Mézard, M., Sausset, F., Sun, Y., and Zdeborová, L. (2011). Statistical physics-based reconstruction in compressed sensing. *arXiv:1109.4424*.

LeCun, Y., Bottou, L., Bengio, Y., and Haffner, P. (1998). Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324.

Lin, M., Chen, Q., and Yan, S. (2013). Network in network. *arXiv:1312.4400*.

Majumdar, S. N. and Schehr, G. (2014). Top eigenvalue of a random matrix: Large deviations and third order phase transition. *Journal of Statistical Mechanics: Theory and Experiment*, 2014(1):P01012.

Mehta, M. L. (2004). *Random matrices*, volume 142. Academic press.

Mezard, M. and Montanari, A. (2009). *Information, physics, and computation*. Oxford University Press.

Mezard, M., Parisi, G., and Virasoo, M. A. (1987). Spin glass theory and beyond.

Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., et al. (2015). Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533.

Mobahi, H. (2016). Training Recurrent Neural Networks by Diffusion. *arXiv:1601.04114*.

Moody, J., Hanson, S., Krogh, A., and Hertz, J. A. (1995). A simple weight decay can improve generalization. *Advances in neural information processing systems*, 4:950–957.

Nalisnick, E., Anandkumar, A., and Smyth, P. (2015). A scale mixture perspective of multiplicative noise in neural networks. *arXiv:1506.03208*.

Neelakantan, A., Vilnis, L., Le, Q. V., Sutskever, I., Kaiser, L., Kurach, K., and Martens, J. (2015). Adding gradient noise improves learning for very deep networks. *arXiv:1511.06807*.

Neyshabur, B. and Panigrahy, R. (2013). Sparse matrix factorization. *arXiv:1311.3315*.

Nguyen, A., Yosinski, J., and Clune, J. (2014). Deep neural networks are easily fooled: High confidence predictions for unrecognizable images. *arXiv:1412.1897*.

Patel, A. B., Nguyen, T., and Baraniuk, R. G. (2015). A Probabilistic Theory of Deep Learning. *arXiv:1504.00641*.

Sagun, L., Guney, V. U., and LeCun, Y. (2014). Explorations on high dimensional landscapes. *arXiv:1412.6615*.

Salakhutdinov, R. and Hinton, G. E. (2009). Deep Boltzmann machines. In *International Conference on Artificial Intelligence and Statistics*, pages 448–455.

Seide, F., Fu, H., Droppo, J., Li, G., and Yu, D. (2014). 1-Bit Stochastic Gradient Descent and its Application to Data-Parallel Distributed Training of Speech DNNs. In *Fifteenth Annual Conference of the International Speech Communication Association*.

Subag, E. (2015). The complexity of spherical p-spin models — A second moment approach. *arXiv:1504.02251*.

Subag, E. and Zeitouni, O. (2015). The extremal process of critical points of the pure $p$-spin spherical spin glass model. *arXiv:1509.03098*.

Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., and Rabinovich, A. (2014). Going deeper with convolutions. *arXiv:1409.4842*.

Talagrand, M. (2003). *Spin glasses: A challenge for mathematicians: Cavity and mean field models*, volume 46. Springer Science & Business Media.

Tao, T. (2012). *Topics in random matrix theory*, volume 132. American Mathematical Soc.

Van der Maaten, L. and Hinton, G. (2008). Visualizing data using t-SNE. *Journal of Machine Learning Research*, 9(2579-2605):85.

Wan, L., Zeiler, M., Zhang, S., Cun, Y. L., and Fergus, R. (2013). Regularization of neural networks using Drop-Connect. In *Proceedings of the 30th International Conference on Machine Learning (ICML-13)*, pages 1058–1066.

Wang, Z., Lu, H., and Liu, H. (2014). Tighten after relax: Minimax-optimal sparse PCA in polynomial time. In *Advances in neural information processing systems*, pages 3383–3391.

## A. Proof of Lemma 1

Consider the model of a deep network discussed in Sec. 2.1 that is given by

$$Y = g'\left(W^{p+1} g'\left(W^{p} \dots g'\left(W^{1} X\right)\dots\right)\right); \tag{S-1}$$

where $g'(x) = \mathbf{1}_{\left\{x \geq \frac{d}{3}\right\}}$ is the thresholding function. We can write (S-1) in terms of the number of "active paths" to get rid of the nonlinearities $g'(\cdot)$, i.e., all paths with non-zero activations from an input $X_i$ to the output $Y$. We then have

$$Y = \sum_{i=1}^{n} \sum_{\gamma \in \Gamma_i} X_i \, W_\gamma; \tag{S-2}$$

where $\Gamma_i$ is the set of all active paths connecting the $i^{th}$ input neuron to the output $Y$. Let $W_\gamma$ be the weight along each path $\gamma$. Now define $\gamma_{i,i_1,\dots,i_p}$ to be an indicator random variable that is 1 iff the path $i \to i_1 \to \dots \to i_p \to Y$, i.e., the one that connects the $i_1^{th}$ neuron in the first hidden layer, $i_2^{th}$ neuron in the second hidden layer and so on, is active. With some abuse of notation, we will use the variable $\gamma$ to denote both the path and the tuple $(i, i_1, i_2, \dots, i_p)$. We now work towards removing the dependency of the input $X$ in (S-2). First split the path into two parts, the first edge on layer 1 and the rest of the path as: $\gamma_{i,i_1,i_2,\dots,i_p} = \gamma_{i,i_1} \; \gamma_{i_1,\dots,i_p}$ (conditional upon $\gamma_{i,i_1} = 1$, the sub-paths are independent). This gives

$$Y = \sum_{i,i_1,i_2,\dots,i_p=1}^{n} \gamma_{i,i_1} \; \gamma_{i_1,\dots i_p} \; X_i \; W_{i,i_1} \; W_{i_1,\dots,i_p}.$$

This expression has correlations between different paths that are introduced by the sum over inputs $X_i$. We approximate these correlations to obtain an expression that resembles the Hamiltonian in (3). Define the net coupling due to the sum over the inputs as

$$Z_i := X_i \; \gamma_{i,i_1} \; W_{i,i_1}, \quad Z := \sum_{i=1}^{n} Z_i.$$

Since $X_i \in \{0,1\}$, we have var $X_i \leq 1/4$ (the variance is maximized for a Bernoulli random variable with probability $1/2$), and $W_{i,i+1}$ is positive or negative with equal probability $d/2n$ while $\gamma_{i,i_1} = 1$ with probability $d/n$ for a fixed $i_1$. Note that E $Z_i = 0$ while E $(Z_i^2) \leq d^2/(4n^2)$. We now use Bernstein's inequality to see that:

$$P\left(\left|\sum_{i=1}^{n} Z_i\right| > t\right) \leq \exp\left(-\frac{t^2/2}{\text{var}\,(Z) + t/3}\right)$$
$$\Rightarrow |Z| \leq n^{-\frac{1}{2}+\frac{1}{p}+\varepsilon} \leq n^{-1/5} \qquad \text{for } p > 5;$$

with high probability for any small $\varepsilon > 0$ (we picked $\varepsilon = 1/5$ above). This also gives var $Z \leq \frac{n^{-2/5}}{4}$. We now get a lower bound on var $Z$ using Chebyshev's inequality, for that we assume that the weight distribution is not all concentrated around zero. In particular, let us assume that $\text{P}\left(\left|W_{ij}^k\right| > t\right) \geq t$ for some small $t = n^{-2/15}$; note that the special form $\text{P}(W > t) \geq t$ is only for convenience. This gives var $W \geq t^3 = n^{-2/5}$ and finally

$$\frac{n^{-2/5}}{4} \geq \text{var}\, Z \geq \text{E}\, X^2 \; n^{-2/5+1/p}. \tag{S-3}$$

We have thus shown that the net coupling due to inputs in our model is bounded, but also has variance of the order $\Theta(n^{-2/5})$.

We now compute the probability of a path $\gamma_{i_1,\ldots,i_p}$ being active. Note that hidden units on the $k^{th}$ layer are uniform, but are not simply Bernoulli with probability $\frac{\rho}{n}\left(\frac{d}{2}\right)^{p-k}$, this is because of the correlations in the layers above. However, for networks that are at most $p=\mathscr{O}(\log_d n)$ deep, we can bound the probability of the $\ell^{th}$ neuron being active on the $k^{th}$ layer (Arora et al., 2013, Proof of Lemma 6):

$$\mathrm{P}\left(h_\ell^k=1\right)\in\left[\frac{3\rho}{4n}\,\left(\frac{d}{2}\right)^{p-k},\frac{5\rho}{4n}\,\left(\frac{d}{2}\right)^{p-k}\right];$$

which gives $P\left(\gamma_{i_1,\ldots,i_p}=1\right)=\Theta(n^{-(p-1)/2})$. The random variable $Z\,\gamma_{i_1,\ldots i_p}$ is zero mean with variance of the order $\Theta(n^{-(p-1)/2-2/5})$ and we therefore replace it with a zero-mean standard Gaussian $K_{i_1,\ldots,i_p}$

$$Y_J\overset{law}{=}\frac{J}{n^{(p-1)/4+1/5}}\sum_{i_1,i_2,\ldots,i_p=1}^{n}K_{i_1,\ldots,i_p}W_{i_1,\ldots,i_p},\tag{S-4}$$

for some constant $J>0$.

Since the entries of weight matrices $W^k$ are iid, the weights along each path $\gamma\in\Gamma_i$, i.e., $W_{i_1,\ldots,i_p}$ are independent and identically distributed. Note that there are at most $\mathscr{O}(m^p)$ distinct values that $W_{i_1,\ldots,i_p}$ can take since each $W_{i_k,i_{k+1}}$ is supported over a set of cardinality $m=\Theta(n)$ in the interval $[-N,\,N]$. For $\sigma_{l_k}\in\mathrm{supp}(W_{i_k,i_{k+1}})$ with $l_k\le n$, we now write $W_{i_1,\ldots,i_p}$ as a sum of weighted indicator variables to get

$$W_{i_1,\ldots,i_p}=\sum_{l_1,\ldots,l_p=1}^{m}\mathbf{1}\left\{W_{i_1,i_2}=\sigma_{l_1},\ldots,W_{i_p,Y}=\sigma_{l_p}\right\}\,\left(\prod_{k=1}^{p}\sigma_{l_k}\right),$$

to get

$$\begin{aligned}Y_J&=\frac{J}{n^{(p-1)/4+1/5}}\sum_{l_1,\ldots,l_p=1}^{m}\sum_{i_1,i_2,\ldots,i_p=1}^{n}K_{i_1,\ldots,i_p}\mathbf{1}\left\{W_{i_1,i_2}=\sigma_{l_1},\ldots,W_{i_p,Y}=\sigma_{l_p}\right\}\,\left(\prod_{k=1}^{p}\sigma_{l_k}\right),\\&\triangleq\frac{J}{n^{(p-1)/4+1/5}}\sum_{l_1,\ldots,l_p=1}^{m}J_{l_1,\ldots,l_p}\left(\prod_{k=1}^{p}\sigma_{l_k}\right);\end{aligned}$$

Note that $J_{l_1,\ldots,l_p}$ is a zero-mean Gaussian with variance $(n/m)^p$ and we can write the scaled output $\widehat{Y}=\frac{n^{-(p-3)/4+1/5}}{\sqrt{m}}\,Y_J$ after a renaming of indices as

$$\widehat{Y}\overset{\text{law}}{=}\frac{J}{m^{(p-1)/2}}\sum_{i_1,i_2,\ldots,i_p=1}^{m}J_{i_1,\ldots,i_p}\,\sigma_{i_1}\ldots\sigma_{i_p}.\tag{S-5}$$

Since each $\sigma_i$ has the same distribution as $W_{ij}^k$, the $p$-product $\sigma_{i_1}\ldots\sigma_{i_p}$ can take at most $\binom{n}{p}$ distinct values for every configuration $\sigma$. Using Stirling's approximation, we can see that this is also of the order $\mathscr{O}(m^p)$ for $p=\mathscr{O}(\log n)$ and $m=\Theta(n)$, i.e., if the discretization of weights grows with the number of neurons we do not lose any representational power by replacing the weights by spins. In the sequel, we will simply set $m=n$ the formulas undergo only minor changes if $m=\Theta(n)$.

Now, the zero-one loss $L$ for $Y^t\sim\mathrm{Ber}(q)$ can be written as

$$\begin{aligned}L=\mathrm{E}_{Y^t}\left|\widehat{Y}(X^t)-Y^t\right|&=q(1-\widehat{Y})+(1-q)\,\widehat{Y}\\&=q+(1-2q)\,\widehat{Y}.\end{aligned}$$

Define $H_{n,p}=(q-L)/(1-2q)$ to see that $-H_{n,p}$ has the same distribution as $\widehat{Y}$.

## B. Gaussian orthogonal ensemble (GOE)

The Hamiltonians considered in this paper are zero-mean Gaussian fields and hence their analysis is tied to the properties of random matrices with Gaussian entries which we describe briefly. Please refer to, for instance, Mehta (2004) or Tao (2012) for an elaborate exposition.

A real symmetric matrix $G \in \mathbb{R}^{n\times n}$ is said to belong to the Gaussian orthogonal ensemble (GOE) if every entry $G_{ij}$ for $i \le j$ is an iid zero-mean Gaussian random variable with $\mathrm{E}\,(G_{ij}^2) = \frac{1+\delta_{ij}}{2n}$. The probability density of GOE matrix is proportional to $\exp\left(-\frac{n}{2}\,\mathrm{tr}\,(G^2)\right)$ and is hence invariant to orthogonal conjugation. The eigenvalues of a real-symmetric GOE matrix are always real and Theorem 2 uses the mean spectral density of eigenvalues of GOE usually characterized as

$$\rho_n(t) = \frac{1}{n}\,\mathrm{E}\sum_{i=1}^{n}\,\delta(t-\lambda_i),$$

where $\lambda_i$ are the eigenvalues. The expected density $\rho_n(t)$ follows the Wigner semi-circle law which states that

$$\lim_{n\to\infty}\mathrm{E}\,\rho_n(t) = \frac{1}{2\pi}\sqrt{2-\lambda^2}\;\mathbf{1}_{\{|\lambda|\le\sqrt{2}\}}. \tag{S-6}$$

The proof of Theorem 2 uses large $n$ asymptotics of the spectral density (Fyodorov, 2013, Eqns. 49-50) around its expected nature given by Wigner's law.

Quite remarkably, the expected number of local minima of an isotropic Gaussian field is obtained by simply replacing the spectral density in Theorem 2 by the density of the largest eigenvalue; this is the subject of Theorem S-1 in Appendix C). In fact, from (S-6) we immediately see that $\mathrm{E}\,\lambda_{\max} \le \sqrt{2}$. More is true, we also have $\mathrm{E}\,\lambda_{\max} = 2$ in the limit, while the fluctuations of $\lambda_{\max}$ around $\sqrt{2}$ for GOE in a region of size $\mathscr{O}(n^{-2/3})$ are given by

$$\sqrt{2}\,n^{2/3}\left(\lambda_{\max}-\sqrt{2}\right) \xrightarrow{d} F_1, \tag{S-7}$$

where $F_1$ is the famous Tracy-Widom distribution. The proof of Theorem S-1 uses asymptotic expressions of the density of $\lambda_{\max}$ for regions beyond $\sqrt{2}\pm\mathscr{O}(n^{-1/3})$ (Majumdar and Schehr, 2014).

## C. Scaling of local minima

The expected number of local minima of the Hamiltonian also undergoes a phase transition based upon the additive perturbations. The following theorem is the corresponding version of Thm. 2 and Thm. 3 that describe phase transitions in the expected number of critical points.

**Theorem S-1** (Fyodorov (2013), Eqn. 60)**.** *For large n, the expected number of local minima of the Hamiltonian in Theorem 2 for $\nu < \nu_c = J\sqrt{p(p-2)}$, i.e., $B > 0$, is given by*

$$E\;crt_0(H) = C(n,B)\;\exp\left\{n\left(\frac{1}{2}\log\frac{1+B}{1-B} - B\right)\right\}. \tag{S-8}$$

*In the transitional regime around $B = 0$, with the scaling $B = -\frac{\kappa}{2}n^{-1/3}$ and $C > 0$, we have*

$$E\;crt_0(H) = \begin{cases} 1 & \text{as } \kappa\to\infty, \\ C\;\exp\left(\frac{\kappa^2}{24} + \frac{4\sqrt{2}\;|\kappa|^{3/2}}{3}\right) & \text{as } \kappa\to-\infty. \end{cases} \tag{S-9}$$

The parameter $\kappa$ interpolates between the two extremes, exponentially many local minima as given by (S-8) for $B > 0$ and constant number of local minima as $B \ll 0$. We again have a transitional regime of $|B| = \mathscr{O}(n^{-1/3})$ where if $\kappa = -2Bn^{1/3}$, we can interpolate smoothly between the totally trivial regime ($\kappa\to\infty$) and the exponential regime ($\kappa\to-\infty$). Roughly, as shown by the following

corollary, we can obtain our desired polynomially many local minima by setting the appropriate value of $\kappa$, viz., $\kappa = \mathcal{O}(\log n)$,

**Corollary S-2.** *If $B = \mathcal{O}\left(n^{-1/3}\ \log n\right)$, the expected number of local minima scales as $\mathcal{O}\left(\exp\left(\log n\right)^2\right)$. In particular, for the spin glass Hamiltonian in (7) and $B$ given by (8), for $p \gg 2$ we require $\nu = Jp\left(1+\kappa n^{-1/3}\right)^{1/2}$.*

Note that for the spin glass model, the number of critical points and local minima is of the same order in all regimes. Moreover, just as we replaced the eigenvalue density by the density of the largest eigenvalue $\lambda_{max}$, we can replace it instead by the probability density of the $k^{\text{th}}$ largest eigenvalue of GOE to get critical points with index $k-1$. A short analysis then shows that there are exponentially many critical points of all indices $k$. Note that the spin glass complexity given by (5) also shows that critical points of all orders are exponentially many.

C.1. **Annealing additive noise with respect to local minima.** We can again construct annealing schemes for additive noise with regularizations using Corollary S-2. For instance, the equivalent expression to (10) for annealing using local minima is $\nu = Jp(1+\kappa n^{-1/3})^{1/2}$ while the equivalent expression for weight decay (cf. (12)) becomes

$$\nu = Jp\left(1+\kappa\frac{\alpha^2}{2n^{1/3}}+\frac{\kappa}{n^{1/3}}\right)^{1/2}.$$

In practice, we choose to anneal based on perturbation magnitude computed for critical points instead of local minima because the former results in reduction of the number of critical points of all indices, in particular, low order saddle points where SGD might not be able to make progress. The latter only guarantees that the number of local minima undergoes a phase transition and does not let us control the number of saddle points.

## D. Quality of perturbed local minima

The main result of this section is Lemma S-7 which says that if the external field is such that if $\nu$ is a constant, the perturbations of the normalized Hamiltonian are $\nu$ in the limit while the perturbations of local minima are at most $\mathcal{O}(n^{-1/3})$. The development here is based on some lemmas in Subag and Zeitouni (2015) which undergo minor modifications for our problem. The basic idea consists of approximating the perturbed Hamiltonian as a quadratic around a critical point of the original Hamiltonian and analyzing the corresponding local minimum.

Let us consider the Hamiltonian

$$\widetilde{H}(\sigma) = H(\sigma) + \frac{1}{\sqrt{n}}\, h^\top \sigma; \tag{S-10}$$

where $h \in \mathbb{R}^n$ is a vector with zero-mean Gaussian iid entries with variance $\nu^2$ and

$$H(\sigma) = n^{-(p-1)/2} \sum_{i_1,i_2,\ldots,i_p} J_{i_1,\ldots,i_p}\ \sigma_{i_1}\ldots\sigma_{i_p}.$$

Note that this is the same model considered in (7) except that we have used a different normalization for the perturbation term. We first get rid off the spherical constraint by re-parameterizing the Hamiltonian around a local minimum $\sigma \in \text{crt}_0(H)$. For $x = (x_1,\ldots,x_{n-1}) \in \mathbb{R}^{n-1}$ with $\|x\| \le 1$ and any $y \in \mathbb{R}^n$ consider the functions

$$P_E(x) = \left(x_1, x_2, \ldots, x_{n-1}, \sqrt{1-\|x\|^2}\right),$$
$$S(y) = \sqrt{n}\, y.$$

Now define

$$f_\sigma(x) \triangleq \frac{1}{\sqrt{n}}\, H \circ \theta_\sigma \circ S \circ P_E(x); \tag{S-11}$$

where $\sigma \in S^{n-1}(1)$ is the normalized spin on the unit-sphere and $\theta_\sigma$ is a rotation that sends the the north pole to $\sigma$. The function $f_\sigma(x)$ is thus the re-parameterization of $H(\sigma)$ around a local minimum $\sigma$, but normalized to have a unit variance. We define $\widetilde{f}_\sigma(x)$ similarly and also define $\overline{f}_\sigma(x)$ for the perturbation term $h^\top \sigma$ to see that (S-10) becomes

$$\sqrt{n}\, \widetilde{f}_\sigma(x) = \sqrt{n} f_\sigma(s) + \overline{f}_\sigma(x).$$

We can now construct the quadratic approximation of $\widetilde{f}_\sigma(x)$ (which we assume to be non-degenerate at local minima, i.e, the Hessian has full rank) to be

$$\sqrt{n}\, \widetilde{f}_{\sigma,\text{approx}}\ (x) \triangleq \sqrt{n}\, f_{\sigma,\text{quad}}(x) + \overline{f}_{\sigma,\text{lin}}(x);$$

where we used the quadratic approximation of the original Hamiltonian and the linear approximation of the perturbation term:

$$\begin{aligned} f_{\sigma,\text{quad}}(x) &= f_\sigma(0) + \langle \nabla f_\sigma(0), x\rangle + \frac{1}{2}\left\langle \nabla^2 f_\sigma(0)x,\ x\right\rangle, \\ \overline{f}_{\sigma,\text{lin}}(x) &= \overline{f}_\sigma(0) + \left\langle \nabla \overline{f}_\sigma(0), x\right\rangle. \end{aligned} \tag{S-12}$$

Let $Y_\sigma$ be the critical point of $\sqrt{n}\, \widetilde{f}_{\sigma,\text{approx}}$ and $V_\sigma$ be its value. We have,

$$Y_\sigma \triangleq -\frac{1}{\sqrt{n}}\, \left(\nabla^2 f_\sigma(0)\right)^{-1}\, \nabla \overline{f}_\sigma(0); \tag{S-13}$$

$$\begin{aligned} \Delta_\sigma &\triangleq \sqrt{n}\, \widetilde{f}_\sigma(Y_\sigma) - \sqrt{n}\, f_\sigma(0) - \overline{f}_\sigma(0), \\ &= -\frac{1}{2\sqrt{n}}\, \left(\nabla \overline{f}_\sigma(0)\right)^\top \left(\nabla^2 f_\sigma(0)\right)^{-1}\, \nabla \overline{f}_\sigma(0); \end{aligned} \tag{S-14}$$

where $Y_\sigma$ is the critical point of $\widetilde{f}_\sigma$. The term $\sqrt{n} f_\sigma(0)$ is the value of the original Hamiltonian at the unperturbed minimum and thus $\Delta_\sigma$ is the perturbation in the critical value of the Hamiltonian due to the perturbations minus the contribution of the perturbation $\overline{f}_\sigma(0)$ itself. We now allude to the following lemma that gives the distribution of the gradient and the Hessian of the Hamiltonian at the critical point, viz., $\nabla f_\sigma(0)$ and $\nabla^2 f_\sigma(0)$.

**Lemma S-3** (Auffinger et al. (2013), Lemma 3.2)**.** *The gradient $\nabla f_\sigma(0)$ is zero-mean Gaussian with variance $p$ and is independent of $\nabla^2 f_\sigma(0)$ and $f_\sigma(0)$. Also, conditional on $f_\sigma(0) = u$, the Hessian $\nabla^2 f_\sigma(0)$ has the distribution*

$$\sqrt{2(n-1)p(p-1)}M - puI; \tag{S-15}$$

*where $M \in \mathbb{R}^{(n-1)\times(n-1)}$ is a GOE matrix and $I \in \mathbb{R}^{(n-1)\times(n-1)}$ is the identity.*

We now compute bounds for $\|Y_\sigma\|$ and $\|\Delta_\sigma\|$ using the following two lemmas but let us describe a few quantities that we need before. It can be shown that the value of a $p$-spin glass Hamiltonian at local minima concentrates around a value

$$m_n = -nE_0 + \frac{\log n}{2\Theta'(-E_0)} - K_0; \tag{S-16}$$

where $E_0$ is the lower energy bound in Fig. 1, $\Theta$ is the complexity of critical points discussed in Fig. 1. In fact, this quantity $m_n$ is the same as the one in Fig. 2d (modulo normalization by $n$), the Hamiltonian is distributed as a negative Gumbel distribution with $m_n$ as the location parameter. The constants $C_0$

and $K_0$ are computed in Subag and Zeitouni (2015, Eqn. 2.6). For a small $\delta \in (0, p(E_0, E_{\rm inf}))$, define $\mathscr{V}(\delta)$ to be the set of real symmetric matrices with eigenvalues in the interval

$$(pE_0 - 2\sqrt{p(p-1)} - \delta, pE_0 + 2\sqrt{p(p-1)} + \delta);$$

and define the set $B_1 = \sqrt{n}\mathscr{V}(\delta)$. The proof of Lemma S-6 hinges upon the observation in Subag (2015) that the additional negative term in the Hessian at a critical point (S-15) does not change the eigenvalue distribution of the GOE much.

**Lemma S-4.** *For any $L > 0$ and*

$$g_3(\sigma) = \Delta_\sigma + \frac{\nu^2}{2\sqrt{n}}\, tr\Big\{\big(\nabla^2 f_\sigma(0)\big)^{-1}\Big\}, \quad B_3 = \Big(-n^{-\frac{1}{2}+\varepsilon}, n^{-\frac{1}{2}+\varepsilon}\Big);$$

*there exists a sequence $e_{\varepsilon,\delta}(n) \to 0$ such that*

$$P\big(g_3 \notin B_3 :\ \nabla^2 f_\sigma(0) = A_{n-1},\ \nabla f_\sigma(0) = 0,\ f_\sigma(0) = u\big) \le e_{\varepsilon,\delta}(n),$$

*for all $u \in n^{-1/2}\,(m_n - L, m_n + L)$ and $A_{n-1} \in \sqrt{n}\ \mathscr{V}(\delta)$.*

*Proof.* First note that $\nabla \overline{f}_\sigma(0) = h$ and the term $\Delta_\sigma$ evaluates to $-\frac{1}{2\sqrt{n}}\ h^\top \big(\nabla^2 f_\sigma(0)\big)^{-1} h$. The gradient of the perturbation $\nabla \overline{f}_\sigma(0) = h$ is independent of the disorder that defines $\nabla^2 f_\sigma(0)$. Thus, conditional on $\nabla^2 f_\sigma(0) = A_{n-1} \in B_1$, we have that $\nabla \overline{f}_\sigma(0)$ is distributed as

$$\nu \sum_{i=1}^{n-1} W_i\, a_i;$$

where $W_i \sim N(0,1)$ are standard Gaussian random variables and $a_i$ are the eigenvectors of $A_{n-1}$. Because $h \sim N(0, \nu^2 I)$, we now see that $g_3(\sigma)$ has the same distribution as

$$\frac{\nu^2}{2\sqrt{n}} \sum_{i=1}^{n-1} (1 - W_i^2)\ \frac{1}{\lambda(A_{n-1})};$$

where $\lambda(A_{n-1})$ are the eigenvalues of $A_{n-1}$ and $W_i \sim N(0,1)$ are standard Gaussian random variables. This is a zero-mean random variable with bounded second moment, indeed

$$\left(\frac{\nu^2 C}{2n}\right)^2 \sum_{i=1}^{n-1} \mathrm{E}\Big\{\big(1 - W_i^2\big)^2\Big\} = 2\left(\frac{\nu^2 C}{2n}\right)^2 (n-1); \tag{S-17}$$

for some constant $C > 0$. The lemma now follows by Chebyshev's inequality. ∎

**Remark S-5.** Note that we had used a different normalization for the magnetic field in (S-10) as compared to (7). Switching back to the original normalization, i.e., replacing $\nu \leftarrow \nu n^{1/2}$, we see from (S-17) that we still have a bounded second moment if $\nu$ is a constant, in particular, if it does not grow with $n$.

**Lemma S-6** (Subag and Zeitouni (2015), Lemma 17)**.** *For $\tau_{\varepsilon,\delta} \to 0$ as $n \to \infty$ slowly enough and $L > 0$*

$$g_4(\sigma) = \frac{\nu^2}{2\sqrt{n}}\, tr\Big\{\big(\nabla^2 f_\sigma(0)\big)^{-1}\Big\}, \quad B_4 = (C_0 - \tau_{\varepsilon,\delta}(n),\ C_0 + \tau_{\varepsilon,\delta}(n));$$

*we have*

$$\lim_{n\to\infty} E\Big\{\big|\sqrt{n}\ \sigma : H(\sigma) \in [m_n - L, m_n + L],\ g_4(\sigma) \notin B_4,\ \nabla^2 f_\sigma(0) \in \sqrt{n}\mathscr{V}(\delta)\big|\Big\} = 0.$$

We can now easily combine the estimates of Lemmas S-4 and S-6 to obtain:

**Lemma S-7.** *The perturbation of the normalized Hamiltonian is $\nu$ in the limit $n \to \infty$ while the perturbed local minima lie within a ball of radius $n^{-\alpha}$ with $\alpha \in (1/4, 1/3)$ centered at the original local minima if $\nu$ is a constant. More precisely for all critical points $\sigma \in crt_0(H)$,*

$$\lim_{n\to\infty} \frac{1}{\sqrt{n}} \left| \widetilde{f}_\sigma(Y_\sigma) - f_\sigma(0) \right| = \nu,$$
$$\|Y_\sigma\| \leq n^{-\alpha}.$$

*Proof.* The above two lemmas together imply that $|g_4(\sigma)| \leq C_0$ with high probability and likewise, $|\Delta_\sigma + g_4(\sigma)| \leq n^{-1/2+\varepsilon}$. Together, this gives that

$$|\Delta_\sigma| \leq C_0 + n^{-1/2+\varepsilon},$$
$$\sqrt{n}\left| \widetilde{f}_\sigma(Y_\sigma) - f_\sigma(0) \right| \leq \|h\| + C_0 + n^{-1/2+\varepsilon}. \tag{S-18}$$

We now set $\nu \leftarrow \nu n^{1/2}$ in (S-10) to match with the normalization used in the main paper to get that the perturbations to the normalized Hamiltonian are

$$\frac{1}{\sqrt{n}} \left| \widetilde{f}_\sigma(Y_\sigma) - f_\sigma(0) \right| \leq \nu + \frac{C_0}{n} + n^{-3/2+\varepsilon}. \tag{S-19}$$

Since the gradient $\nabla f_\sigma(0)$ is independent of $\nabla^2 f_\sigma(0)$ and $f_\sigma(0) = u$, we have the following stochastic domination

$$\|Y_\sigma\| \leq \frac{c}{n}\, \|h\|;$$

for some constant $c > 0$. The conditional probability in Lemma S-4 is now computed to be

$$\mathrm{P}\left( Q_{n-1} \geq \frac{n^{\frac{1}{2}-\alpha}}{c\nu} \right); \tag{S-20}$$

where $Q_{n-1}$ is a standard Chi random variable with $n-1$ degrees of freedom, which goes to zero if $\nu$ is a constant that goes not grow with $n$. In particular, we have that $\|Y_\sigma\| \leq n^{-a}$. ∎

## E. Additive noise with Drop-Out and Drop-Connect

Drop-Out (Hinton et al., 2012b) is a form of regularization typically used for fully connected layers where each output is forwarded with a fixed probability $q$ and set to zero otherwise. Numerous experiments have shown that drop-out leads to a better generalization error for deep networks by forcing each layer to explain the layer above using as few features as possible. Drop-Connect (Wan et al., 2013) is very similar to drop-out except that one sets the individual weights to zero with probability $1-q$. This induces sparsity in the weight matrix at each layer. We model drop-connect by introducing a Bernoulli random variable $\delta_{i_1} \dots \delta_{i_p}$ that multiplies the disorder $J_{i_1,\dots,i_p}$ to give

$$-\widetilde{H}_{\text{drop-connect}}(\sigma) = \sum_{i_1,i_2,\dots,i_p} \delta_{i_1} \dots \delta_{i_p}\, J_{i_1,\dots,i_p} \sigma_{i_1} \dots \sigma_{i_p} + \sum_i h_i \sigma_i; \tag{S-21}$$

where as before $J_{i_1,\dots,i_p}$ is zero-mean Gaussian with variance $J^2 n^{-(p-1)}$. Such a Hamiltonian is known as a "diluted" spin glass. Since $\delta_{i_1} \dots \delta_{i_p} = 0$ is like setting the entire path $(i_1, \dots, i_p)$ to zero, we set $\delta_{i_1} \dots \delta_{i_p} \sim \mathrm{Ber}(q^p)$. Similar to the proof of Lemma 1, we now replace $\delta_{i_1} \dots \delta_{i_p} J_{i_1,\dots,i_p}$ by a Gaussian random variable to put it into the framework of Sec. 3. If we denote $c \triangleq q^p$, this amounts to scaling the constant $J$ by $\sqrt{c(1-c)}$ which gives var $J_{i_1,\dots,i_p} = n^{-(p-1)}\, c(1-c)J^2$. For the perturbation term to result in polynomially many critical points, we now require

$$\nu = J\, p\sqrt{c(1-c)}\, \left(1 + \frac{2\tau}{n}\right)^{1/2}. \tag{S-22}$$

Quite intuitively, since $c(1-c) < 1/2$, drop-connect requires that the magnitude of the perturbation term be reduced by an appropriate factor. As a minor aside, following the idea that drop-connect sets entire "paths" from data layer to loss layer to zero we should pick the drop-out probability such that $q^p$ is a constant, i.e., $q = c^{1/p}$.

To model drop-out, we multiply each spin in (7) by $\delta_i \sim \text{Ber}(q)$ to get

$$-\widetilde{H}_{\text{drop-out}}(\sigma) = \sum_{i_1,i_2,\ldots,i_p} J_{i_1,\ldots,i_p} \prod_{k=1}^{p} \delta_{i_k}\, \sigma_{i_k} + \sum_i h_i \sigma_i. \tag{S-23}$$

For our model $J_{i_1,\ldots,i_p}$ is such that each neuron $i_k$ is connected to $i_{k+1}$ in the next layer uniformly randomly. This enables us to interpret the factor $\prod_k \delta_{i_k}$ as a path $i_1 \to \ldots \to i_p$ similar to (S-21) and gives the same expression for the perturbation as (S-22).